

(12) UK Patent Application (19) GB (11) 2 336 022 (13) A

(43) Date of A Publication 06.10.1999

(21) Application No 9906881.9

(22) Date of Filing 25.03.1999

(30) Priority Data

(31) 09055078

(32) 03.04.1998

(33) US

(71) Applicant(s)

Discreet Logic Inc
(Incorporated in Canada - Quebec)
5505 St-Laurent Blvd, Suite 5200, Montreal,
Quebec H2T 1S6, Canada

(72) Inventor(s)

Dale Matthew Weaver

(74) Agent and/or Address for Service

Atkinson & Co
First Floor, Unit A, The Technology Park,
60 Shirland Lane, SHEFFIELD, S9 3PA,
United Kingdom

(51) INT CL⁶

G11B 27/028

(52) UK CL (Edition Q)

G5R RB81

(56) Documents Cited

GB 2289558 A

US 5726717 A

(58) Field of Search

UK CL (Edition Q) G5R RB81

INT CL⁶ G11B 27/028 27/034

Online: EPODOC; JAPIO; WPI

Best Available Copy

(54) Abstract Title

Edit processing audio-visual data

(57) Audio data and visual data are processed by providing storage devices (211, 212) for storing digital samples which are then read in response to edit events defined by a timeline (352, 353). Active audio tracks are processed in real time and an event, such as a cross-fade, defined by a single track that requires two interacting tracks of material, is identified. The location of events within available tracks is re-arranged so as to reduce the total number of active tracks required during playback. In effect, events are transferred to blank regions of existing tracks, thereby allowing both of the interacting tracks of cross-fades to be read without requiring additional track capacity during playback. In this way, an event defined by two interacting tracks is processed without requiring additional processing capacity.




Figure 3

GB 2 336 022 A




Figure 1




Figure 2




Figure 3




Figure 4

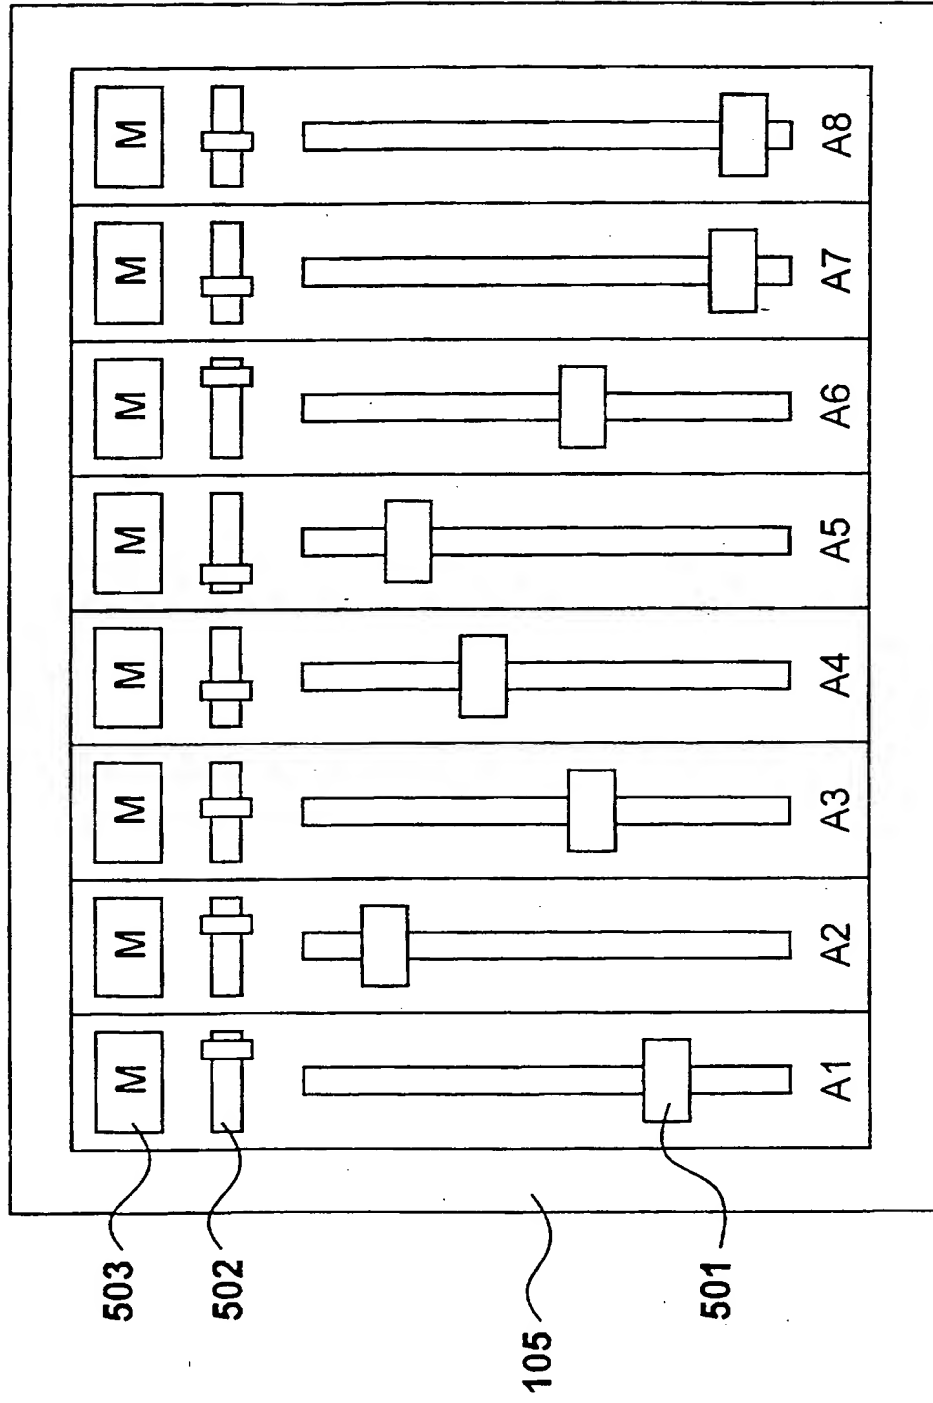

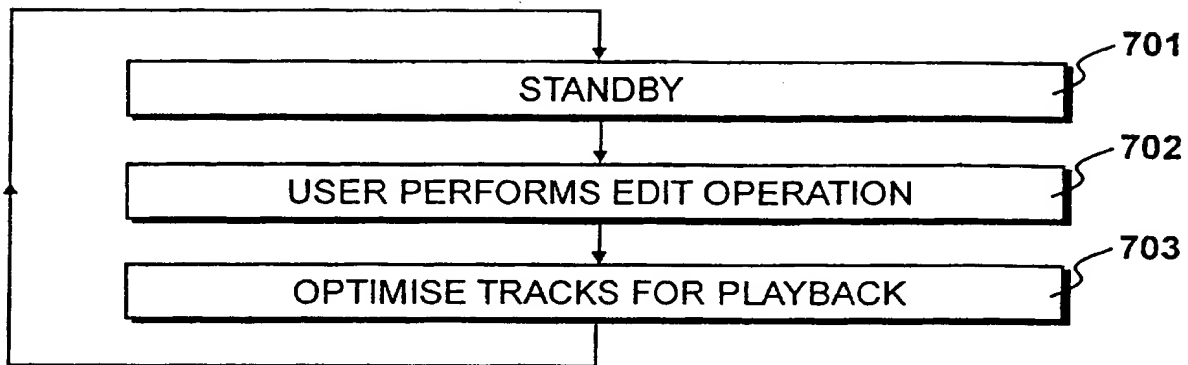



Figure 5

6/11

*Figure 6*

7/11

*Figure 7A**Figure 7B*




Figure 8A

9/11

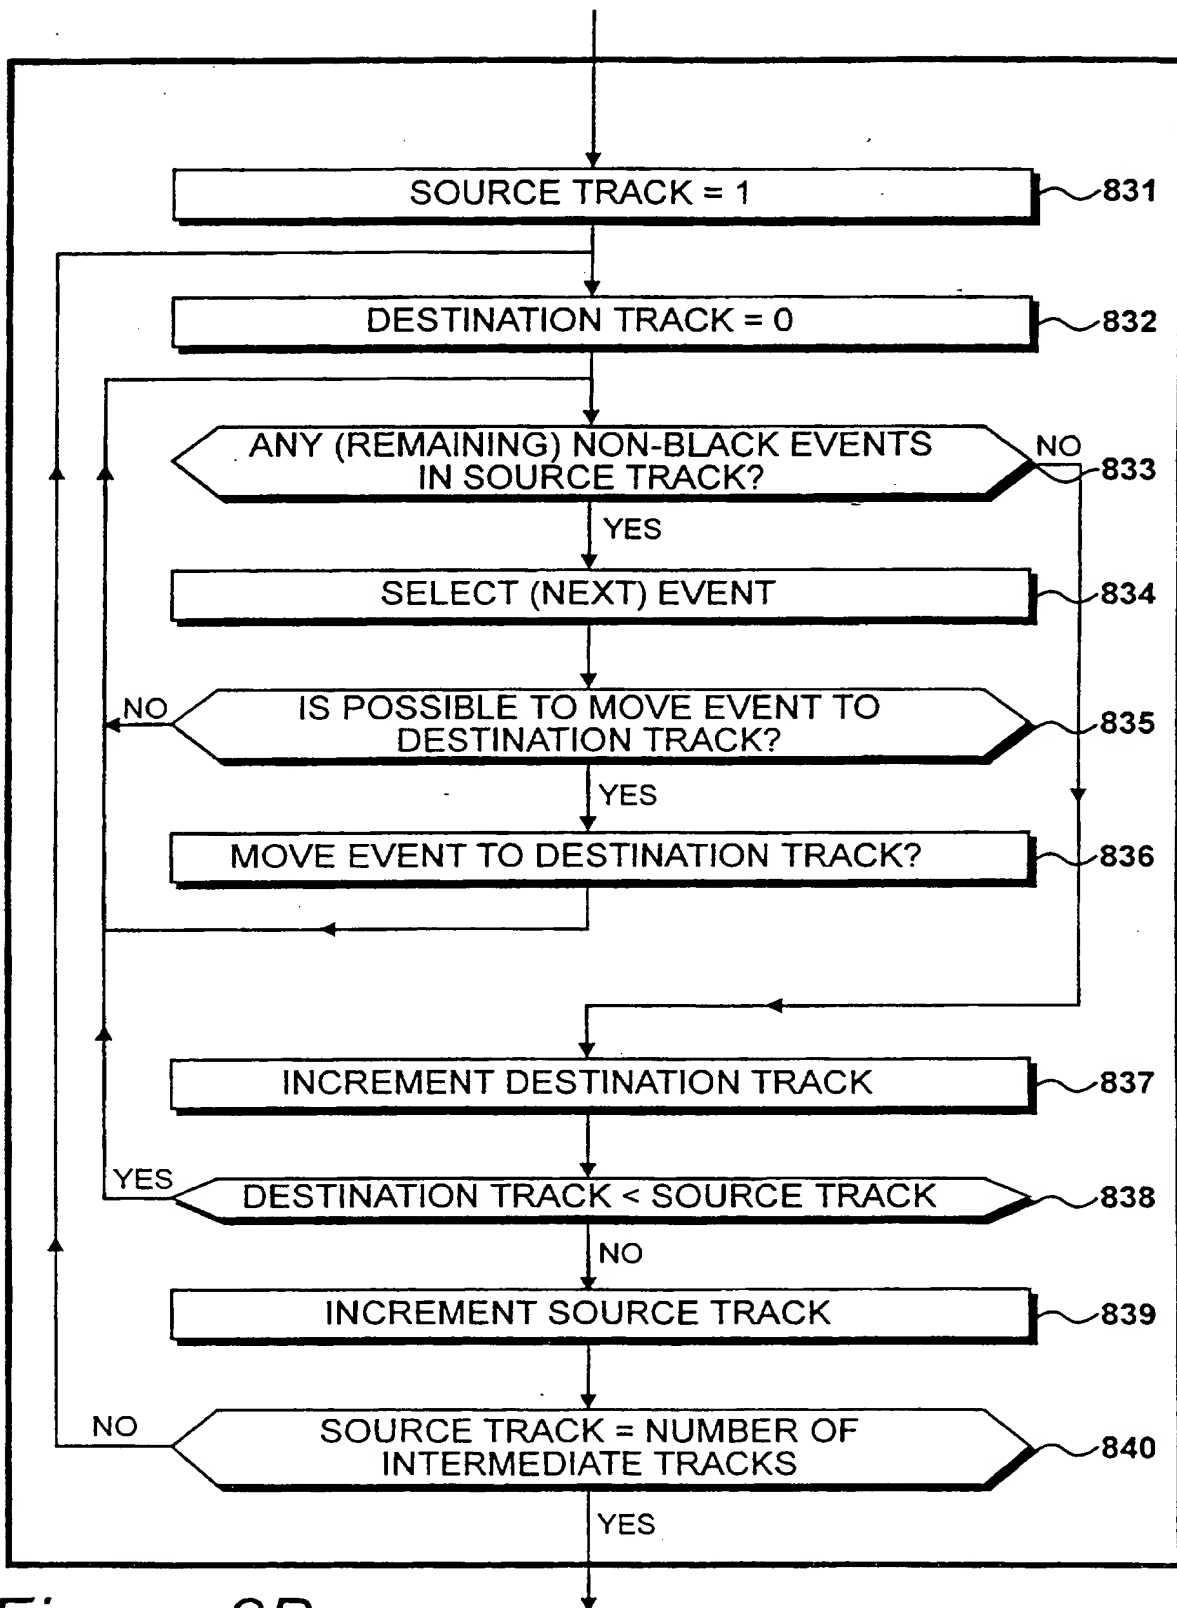


Figure 8B


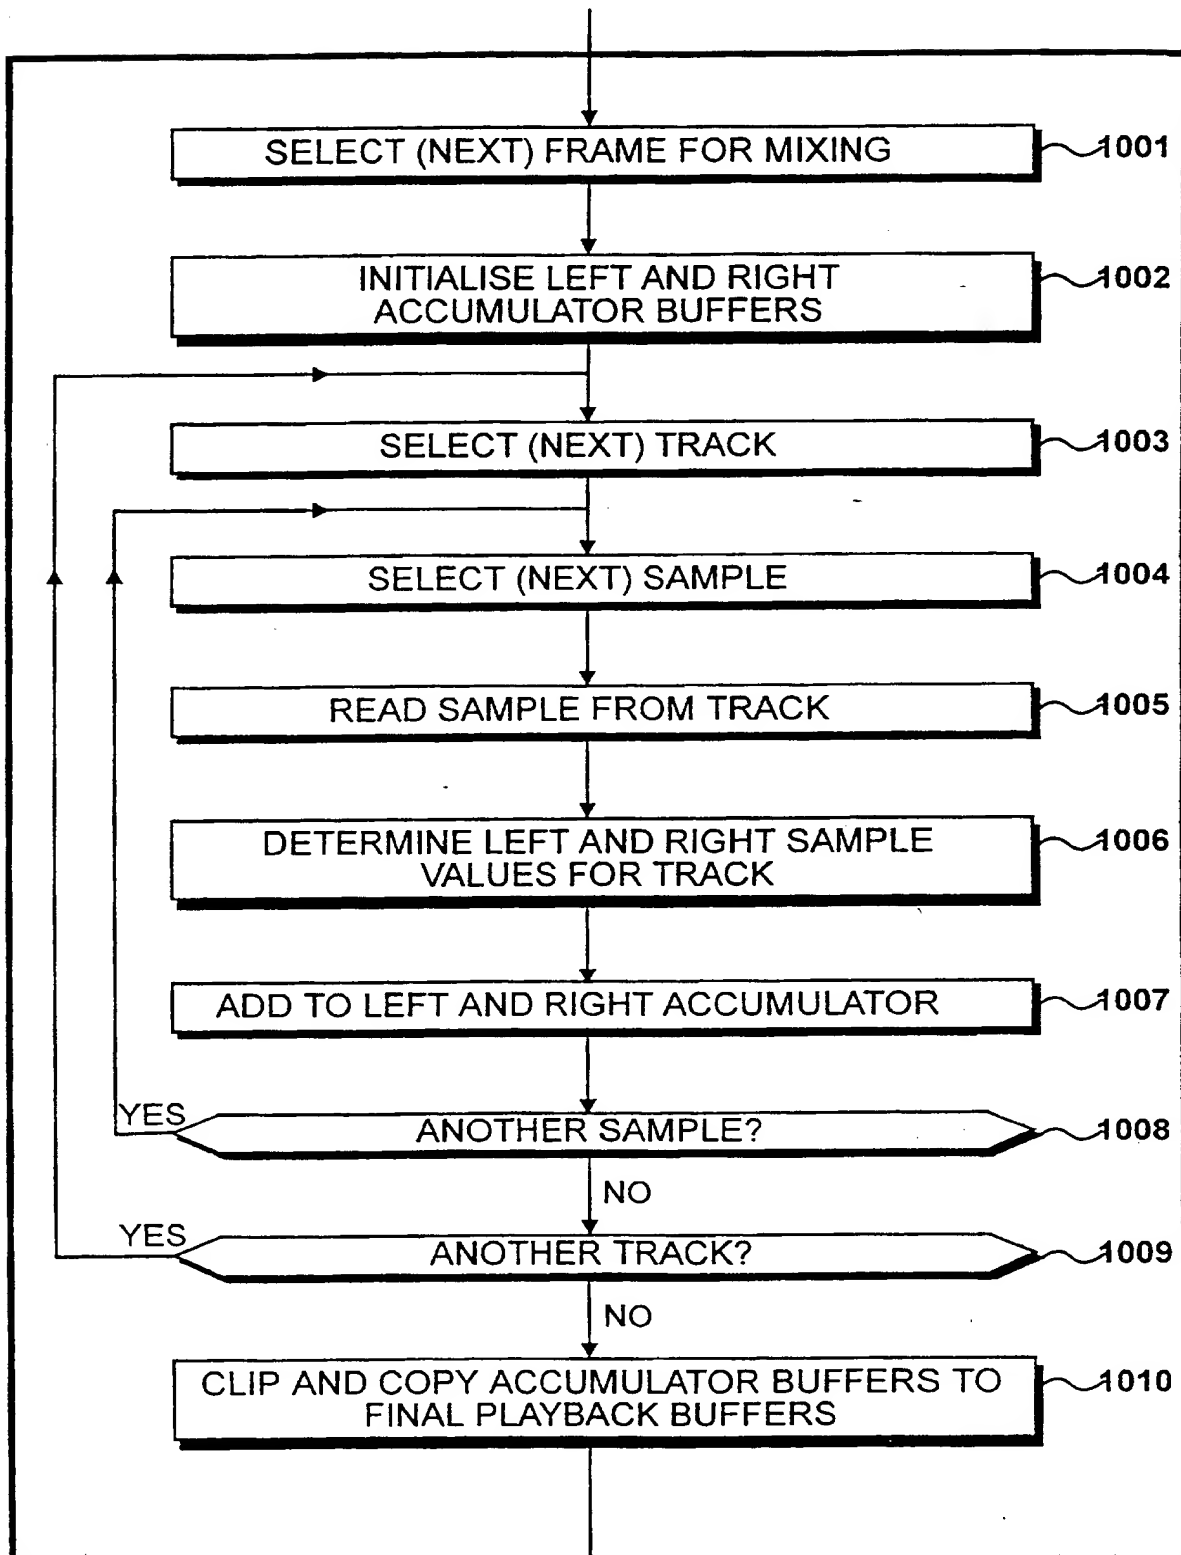


Figure 9

11/11

*Figure 10*

2336022

1

Processing Audio-Visual Data

5 The present invention relates to processing audio-visual data in which digital samples are read from storage media in response to edit events defined in a time line.

10 Traditional video editing involves the copying of video material from source tape onto an edited tape. Sophisticated tape editing equipment is required and the process can be relatively time consuming, given that it is necessary to configure the equipment in order for the video material to be transferred correctly. Furthermore, editing of this type leads to image degradation therefore the number of layers that may be introduced for compositing is limited.

15 In order to optimise expensive on-line editing equipment, off-line editing systems are known in which compressed video images are manipulated rapidly, by accessing image data in a substantially random fashion from magnetic disc storage devices. Given that it is not necessary to spool linearly through lengths of video tape in order to perform editing of this type, the editing process has generally become known as "non-linear editing". Initially, systems of this type would generate edit decision lists
20 such that the on-line editing process then consists of performing edits once in response to an edit decision list. However, the edit decision list itself could be created in a highly interactive environment allowing many potential edits to be considered before a final list is produced.

25 The advantages of non-linear editing have been appreciated and high-end systems are known, such as that licensed by the present assignee under the Trade Mark "FIRE" in which full bandwidth signals are manipulated at full definition, without compression.

In a high-end system, it is possible to specify hardware requirements in order to provide a required level of functionality. Thus, systems tend to

be designed to achieve a specified level of service; being tailored to suit a user's particular demands. However, as the power of processing systems has increased, along with an increase in data storage volumes and access speeds, it has become increasingly possible to provide sophisticated on-line non-linear editing facilities for more general purpose platforms. However, when working with such platforms the extent to which hardware facilities may be enhanced in order to provide particular functionality is more limited. Consequently, there is a greater emphasis towards providing enhanced functionality by making optimum use of the processing capacities available.

According to a first aspect of the present invention, there is provided editing apparatus, including storage means configured to store digital samples; display means configured to display symbolic representations of edit events within tracks; and processing means configured to identify event locations and to move portions of edit events to alternative tracks so as to enhance processing performance.

In a preferred embodiment, the processing means is configured to identify an event defined on a single track but requiring two interactive tracks and to transfer material to one of said interacting tracks to a blank region of another track, thereby allowing both of said interacting tracks to be played without allocating additional track resource.

According to a second aspect of the present invention, there is provided a method of processing audio-visual data in which digital samples are read from storage media in response to edit events identified symbolically within tracks, wherein an improved event location is identified on a different track, and said identified event is moved to said improved location so as to reduce the overall processing requirement.

The invention will now be described by way of example only, with reference to the accompany drawings of which:

Figure 1 shows a non-linear digital editing suite, having a processing system, monitors and recording equipment;

Figure 2 details the processing system shown in *Figure 1*, including a memory device for data storage;

5 *Figure 3* shows a typical time-line display displayed on a VDU of the type shown in *Figure 1*;

Figure 4 details audio tracks shown in *Figure 3*;

Figure 5 illustrates an attribute window displayed on one of the monitors shown in *Figure 1*;

10 *Figure 6* illustrates the arrangement of data contained within the memory device shown in *Figure 2*;

Figure 7A illustrates the editing of audio data by the system shown in *Figure 1*, including a step for the optimisation of tracks for audio playback;

15 *Figure 7B* illustrates the playing of audio data optimised in *Figure 7A*, including a step of mixing audio data;

Figure 8A details the track playback optimisation process identified in *Figure 7*, including a step of optimising intermediate tracks;

Figure 8B details the step of optimising intermediate tracks identified in *Figure 8A*;

20 *Figure 9* details the effect of the optimisation procedures shown in *Figure 8* when applied to the audio data tracks shown in *Figure 3*; and,

Figure 10 details the step mixing audio data identified in *Figure 7*.

25 A non-linear editing suite is shown in *Figure 1* in which a processing system 101 receives manual input commands from a keyboard 102 and a mouse 103. A visual output interface is provided to an operator by means of a first visual display unit (VDU) 104 and second similar VDU 105. Broadcast-quality video images are supplied to a television type monitor 106 and stereo audio signals, in the form of a left audio signal and a right audio signal, are supplied to a left audio speaker 107 and to a right audio

speaker 108 respectively.

Video source material is supplied to the processing system 101 from a high quality tape recorder 109 and edited material may be written back to said tape recorder. Recorded audio material is supplied from system 101 to an audio mixing console 110 from which independent signals may be supplied to the speakers 107 and 108, for monitoring audio at the suite, and for supplying audio signals for recording on the video tape recorder 109. Operating instructions executable by the processing system 101 are received by means of a computer-readable medium such as a CD ROM 111 receivable within a CD ROM player 112.

Processing system 101 is detailed in *Figure 2* and may be summarised as a dual "Pentium Pro" machine. A first processing unit 201 and a second processing unit 202 are interfaced to a PCI bus 203. In the example shown, the processors 201 and 202 are clocked at 233 MHz and these devices communicate directly with an internal memory 204 of one hundred and twenty-eight megabytes, over a high bandwidth direct address and data bus, thereby avoiding the need to communicate over the PCI bus during processing except when other peripherals are being addressed. In addition, permanent data storage is provided by a host disc system 205 of four gigabytes, from which operating instructions for the processors may be loaded to memory 204, along with user generated data and other information. In addition to the host environment, Small Computer System Interface (SCSI) controllers 206, serial interfaces 207, an audio-visual/subsystem 208 and desktop display cards 209 are also connected to the PCI bus 203.

SCSI controllers 206 interface video storage devices 211 and audio storage devices 212. In the present embodiment, these storage devices are contained within the main system housing shown in *Figure 1* although, in alternative configurations, these devices may be housed externally.

Video storage devices **211** are configured to store compressed video data and typically said video data is striped across four, nine-Giga-byte drives. A similar arrangement is provided for the audio storage devices **212** which typically consist of four, four-Giga-byte drives again configured as a striped array. Sufficient bandwidth is provided, in terms of the video storage devices **211** and the SCSI controllers **206**, to allow two video streams of data to flow over the PCI bus **203** in real time. Although the video data is compressed, preferably using conventional JPEG procedures, the data volume of video material is still relatively large compared to the data volume of the audio material. Thus, the audio storage devices **212** in combination with SCSI controllers **206** provide sufficient bandwidth for in excess of one hundred audio channels to be conveyed over the PCI bus **203** in real time.

Serial interfaces **207** interface with control devices **102** and **103** etc. via an input/output port **213**, in addition to providing control instructions for video tape recorder **109** via a video interface port **214**. The video interface port **214** also receives component video material from the audio-visual subsystem **208**.

The audio-visual subsystem **208** may include a Truevision Targa 2000 RTX board configured to code and decode between uncompressed video and JPEG compressed video at variable compression rates. A limited degree of signal processing is provided by subsystem **208**, under the control of the CPUs **201/202** and audio output signals, in the form of a left channel and a right channel, are supplied to an audio output port **215**. Television monitor **106** receives luminance and chrominance signals from subsystem **208** via a video monitor interface **216** and a composite video signal from subsystem **208** is supplied to the desktop display subsystem **209**, via link **217**.

The system operates under the operating system "Windows NT" and is preferably configured under "NT 4.0". Desktop display **209** includes two

VDU driver cards operating in a dual monitor configuration, thereby making the resources of both cards available to the operating system, such that they are perceived as a single large desktop configured with 2048 by 768 pixels. The desktop drivers support video overlay, therefore video sequences from the audio-visual subsystem 208 may be included with the VDU displays in response to receiving the composite signal via link 217. Thus, VDU 104 is connected to VDU interface 218 with VDU 105 being connected to interface 219. However, in operation, the VDUs provide a common desktop window, allowing application windows to be arranged on the desktop in accordance with user preferences.

The editing suite shown in *Figure 1* facilitates timeline editing by displaying timelines on monitor 104, as shown in *Figure 3*. Frame numbers 301 to 322 are shown at the top of the display representing timecode for output frames. Individual output frames may be viewed and a particular output frame may be selected by means of a vertical position line 351. Thus, as output images are being displayed, position line 351 traverses across the image from left to right. In its present position, position line 351 is identifying frame 309.

For the purposes of this example, it is assumed that video material and audio material may be processed in real time when accessing source material from a total of two video sources in combination with source material from a total of six audio sources. Many more than six audio sources could be made active but when more than six audio sources or tracks are made active for playback purposes, real time operation cannot be guaranteed. In a video source track, such as track V1 or track V2, source material is identified by a timeline and reference to the selected source material is included within the timeline. Thus, source material is identified in video timeline V1 in which a cut occurs after frame 316 from video source material 352 to source material 353.

During the playing of video source material **352**, audio material is being played from audio tracks A1,A2,A3,A4 and A5. After the transition, such that video source material is received from source **353**, audio material is received from tracks A1,A2 and A6.

5 A cut, as illustrated with respect to video track V1, is relatively easy to achieve given that from frame **317** onwards video material is read from source **353** instead of being read from source **352**. Edit points are selected in the source material and source timecodes are stored such that the required material is read from its correct position when required in the
10 output stream. In addition to cuts, it is also possible to define a wipe or a dissolve such that, over a transition period, material is derived from two sources with gradual mixing occurring from one source to the other.

 In order to provide a coherent editing environment, effects similar to wipes and dissolves may be specified in the audio environment. In a video
15 dissolve, one image is gradually replaced by another at each location throughout the image. Thus, a similar effect may be achieved with the audio signals by gradually decreasing the volume of one source while simultaneously increasing the output volume of another. In audio systems, such a procedure is usually referred to as a cross fade, given that the first
20 source is being faded down while the second source is being faded up.

 In the example shown in *Figure 3*, an audio dissolve or cross-fade has been specified for audio track A1 by means of an audio dissolve icon **354**. From an operators point of view, audio source material **355** is placed into audio track A1 up to an including output frame **316**. At frame **317** track
25 A1 cuts to audio source **356**, thereby creating a similar transition to that provided for video track V1. Transition effects are then selected and an audio dissolve icon **354** is dragged and dropped at the cut transition between sources **355** and **356**.

 Thus, the audio output from track A1 still consists of source **355**

being played followed by source 356 being played. However, there is no longer an abrupt cut from one source to the other. Instead, source 355 starts to be faded down from frame 313 with source 356 being faded up from this position. Thus, over frames 313 to 320, audio material is derived from both source 355 and from source 356. This requires two input audio streams to be processed and it is not possible for both of these audio streams to be physically supplied to the processing system via audio track A1. In reality, source 355 must be extended from its notional cut position to the end of frame 320, to provide source material for the fade out. At the start of frame 313 source material is required for source 356 and, in order to provide this source material to the processing system, it must be made available by another audio track.

The audio tracks shown in *Figure 3* are detailed in *Figure 4*, in which account has been taken of the fact that, during the audio dissolve or cross fade 354, two audio sources are required to be active in order to satisfy the effect specified in audio track A1. During the playback of audio track A1, two audio sources are required over frames 313 to frame 320. This requirement is satisfied if material is replayed via an additional active track. Thus, as shown in *Figure 4*, the requirement for input material 356 may be transferred from audio track A1 to new audio track A7. Thus, from frame 313 to frame 320, source material 355 is supplied via audio track A1 with source material 356 being to supplied via audio track A7.

After frame 320, audio source 356 could continue to be supplied via audio track A7, as indicated by outline 401. However, the system only guarantees the ability to play back six audio tracks in real time, therefore it is preferable for source material 356 to be replayed via audio track A1 from frame 321 onwards, so that audio track A7 may be muted. Thus, it is possible for audio track A7 to be rendered active only for the period during which it is required, whereafter the track is automatically muted so as to

reduce processing burden. However, over the period of output frames 313 to 320, it is necessary to receive and process audio signals from seven audio tracks; a situation which is likely to result in system degradation.

In addition to representing time line edits as shown in *Figure 3*, an attribute window may be selected, and displayed on monitor 105, as shown in *Figure 5*. The attribute window allows attributes to be defined for each of the audio tracks. Sliders are presented to an operator on monitor 105 and an operator may adjust these sliders by operation of the mouse 103. A first slider 501 allows the overall volume level of audio-track A1 to be adjusted. In addition, a panning slider 502 allows the pan, i.e. stereo position, of audio channel A1 to be selected. The audio channel is also provided with a mute button 503 such that, when selected, audio channel A1 is muted, thereby reducing processing burden.

Thus, for each of the audio tracks it is possible to define attribute data in terms of volume and pan on a frame by frame basis. Alternatively, if required, attributes for volume and pan may be stored at sub-frame definition and, in the ultimate case, volume and pan values could be stored for each individual audio sample. However, in the preferred embodiment, audio attributes, specifically volume and pan values, are specified on a frame by frame basis for each frame period of each audio track. During real time playback, sub-frame volume and pan values are computed for each one sixteenth of a frame using linear interpolation between the attribute values specified by the user at the neighbouring frame boundaries.

Audio samples will generate a click if attenuated too abruptly. Thus, when a user specifies a cut at an event in or out point, the system can be set up so as to automatically ramp up or down to or from the maximum amplitude specified for the track. The user can specify the duration of the ramp in or out, in terms of a percentage of frame length, from zero to one hundred percent. If set to zero, there will be no ramp effect. This situation is

suitable for conditions where the track contains only very low level background noise, and no resulting click would be audible. When set to one hundred percent, the ramp is at its slowest, and the audio signal is introduced or faded out slowly enough that no clicks or aliasing artefacts are audible. The sample amplitude factor, that changes in response to the ramp effect being defined in this way, is therefore computed on a sample by sample basis, so that stepping is avoided, and a smooth response is obtained during playback. It should be appreciated that more audio tracks may be processed if said tracks are processed at lower bandwidths. However, in the preferred embodiment, a sampling rate of 48 kHz is used with sixteen bits being allocated to each left and to each right audio sample.

In order to process each audio track, it is necessary to calculate left channel and right channel contributions by performing a multiplication upon each audio sample. Thus, processing units **201** and **202** are required to manipulate samples on a sample by sample basis in order to generate the output data. To facilitate this process, audio data from the audio data store **212** is written to relatively large buffers within memory **204**, as shown in *Figure 6*. Within the overall system memory **204**, instructions for the operating system are stored at location **601**. Similarly, application instructions are stored at locations **602** with the locations above **602** being available for the storing and buffering of application data.

Locations **603** provide a first audio buffer with locations **604** providing a second audio buffer. In this way, double buffering is facilitated such that one of said buffers may receive data from the audio store **212** while the other buffer is providing data to be mixed. Thus, data may be written to a buffer efficiently as a substantially constant stream from the storage devices **212** whereafter the data may be accessed randomly from the other buffer.

Mixed data is written to a playback buffer in storage location **605** and

this data is then read to provide a frames-worth of mixed data to the audio-visual subsystem **208**.

The transfer of data over the PCI bus **203** is effectively interrupt driven and CPU **201** is responsible for bus mastering and for controlling interrupt priorities. However, the operations performed may be visualised in
5 a more systematic way, as illustrated in *Figures 7A and 7B*.

At step **701** operations effectively remain in a standby mode until the user requests an operation to be performed. At step **702** the user has performed an edit operation that results in a change being made to the
10 event structure of the tracks, and therefore possibly resulting in a change in the way that the tracks can be optimised. As a result of this change, at step **703**, the tracks are optimised for playback. The audio-visual data stored on storage devices **211** and **212** is in the form of digital samples which are read in response to edit events defined by the timeline shown in *Figure 3*. A
15 plurality of active tracks may be processed in real time and the optimisation process is performed so as to mitigate the effect of reading data which has been created as a single track but in actual fact requires two tracks of material to be read. An event of this type, such as a dissolve, wipe or similar audio transition, is optimised by optimisation process **703**.
20 Conventionally, upon detecting such a situation, it would be necessary to activate another track and to require the reading of data from the new active track which, in some circumstances, may degrade the operation of the system. However, in the present embodiment, material for one of the interacting tracks is associated with a blank region of another active track
25 thereby allowing both of the interacting tracks to be read without requiring an additional active track. Thus, optimising step **703** consists of optimising track events for conditions such as audio dissolves and reallocating a portion of track material to a blank portion of an already active track, so as not to require a non-active track to be activated. After optimisation step

703, control is directed back to step 701, where the process enters a standby mode of operation until the next event occurs. The sequence of operations illustrated in *Figure 7A* does not attempt to fully describe the processes performed by the processing system shown in *Figure 2*, as there are many processes being performed in accordance with multitasking protocols under the control of the operating system. What *Figure 7A* illustrates, is that optimisation is performed when the user defines a change in data in such a way that it then becomes necessary for a new optimisation to be performed.

As a result of optimisation 703 shown in *Figure 7A*, the processes shown in *Figure 7B* may be effected with greater efficiency. The transfer of video data from video store 211 to the audio/video subsystem 208 is performed under the control of CPU 201, and the majority of the actual video processing is performed by the subsystem 208. The procedures identified in *Figure 7B* are therefore directed towards the processing of the audio data. At step 710, the process is in a standby mode. At Step 711, the user has requested a playback to begin. Optimisation has already been carried out at step 702 in preparation for playback, which then proceeds at step 711. At step 711 data for playback is identified. This includes data already held in local memory 204, but typically will also require access to long term memory storage, such as the hard discs 211 and 212. At step 711, therefore, preparations are made to ensure that data is available with sufficient speed to keep up with demands for real time playback. At step 712, a batch of frames is identified for output. The number of frames selected depends on a number of factors, including overall processor usage, and the size of free memory. Step 712 attempts to ensure that efficient use is made of the available processing resources. For example, if this was not done, it is possible that an attempt would be made to allocate memory resources for data beyond those which are currently available,

resulting in a time consuming delay while resources are adjusted. Avoidance of these types of delays is crucial to obtaining real time performance.

At step 713 new data is read from the audio disc 212 and written to
5 buffer number one. At step 714 data is read from buffer number two and
mixed so as to write stereo data to the playback buffer 605. At step 715
buffer duties for audio buffer number one and audio buffer number two are
exchanged and at step 716 the mixed audio data from the playback buffer
605 is supplied over the PCI bus 203 to the audio-video subsystem 208. A
10 question is asked at step 717 as to whether another batch of frames is to
be processed and when answered in the affirmative control is returned to
step 712. Eventually, the playback will have been completed and the
system will return to its standby condition at step 710.

Without optimisation process 703, the dissolve 354 shown in *Figure*
15 3 would result in audio channels being processed as illustrated in *Figure* 4.
The dissolve 354 may be considered as a portion of the playback during
which a tail of source material 355 is being processed in combination with a
head of material 356. In order to allocate the material onto individual audio
tracks, it would be possible to move the tail of track 355 or to move the
20 head of track 354. By convention, in the present embodiment, tails are
retained in their specified tracks and heads are transferred.

Step 702 for the optimisation of tracks for playback is shown in
Figure 8A. At step 801 a question is asked as to whether more than one
audio track has been selected for playback. If answered in the negative, no
25 optimisation is necessary. Alternatively, if optimisation is necessary, control
is directed to step 802. At step 802, a number of intermediate tracks are
generated. The intermediate tracks are never noticed by the user, but are
generated for efficient use of processing resources when audio tracks are
being played. For a number N of audio tracks defined by the user, 2N

intermediate tracks are required, thus an additional N tracks are created in addition to the original user-specified tracks. At step 803 a track is selected from the first N intermediate tracks. On the first pass through the processes illustrated in *Figure 8A*, the track selected will be track 0. On the next pass, the track selected will be track 1, and so on, up to track N-1. The selected track number is designated as track M. At step 804 a question is asked as to whether the currently selected track contains a dissolve, or in other words, does the track contain data that in reality requires a pair of tracks during playback. A dissolve has this requirement, as two sources of audio data are combined because it requires two audio tracks to be made available simultaneously. If the question asked at step 804 is answered in the negative, control is directed to step 806. Alternatively, control is directed to step 805.

At step 805, the presence of dissolve material has been identified. The dissolve material comprises two parts: a head portion and a tail portion. The head portion is the material on the track leading up to and including material contributed during the dissolve, and the tail portion is the remaining material on the track including material contributing to the dissolve. At step 805, the head portion is placed on track N+M, and the remaining material, including the tail portion, is placed on track M. In step 805, this is performed for each instance of a dissolve that is encountered on the track. In every case, material is added with attribute data to show which user track the material came from, whether it came from the head or tail of a transition, or neither, as well as other data needed for real time playback.

After step 805, control is directed to step 806, where a question is asked as to whether another track is to be considered. This condition is satisfied when M is less than N, in which case control is directed back to step 803. Alternatively, if all of the original source tracks have been considered, control is thereafter directed to step 807. At step 807 the

arrangement of material on the intermediate tracks is optimised, resulting in a reduction in the number of intermediate tracks required during playback. At step 808, empty tracks, resulting from the optimising process at step 807, are deleted. This reduces the processing requirements for playback.

5 At step 809 a question is asked as to whether the total number of remaining tracks is too high to guarantee real time playback. In the present embodiment, this number is six. In an alternative embodiment, the number could be variable, depending upon the other tasks required to be performed by the processing system, for example, the number of video channels being

10 played back. If the number of audio channels is too high, control is directed to step 810, where an audio overload condition is set. Alternatively, it is known that the processor can provide sufficient power to play back the audio tracks in real time, and it is unnecessary to set the audio overload condition. This completes the operations for optimising tracks for playback.

15 The process for optimising intermediate tracks, indicated at step 805 in *Figure 8A* is detailed in *Figure 8B*. At step 831 the source track is selected initially as being track one. At step 832, the destination track is selected initially as being track zero. For the purposes of this flow chart, tracks are considered as being numbered from zero to $2N-1$, where $2N$ is

20 the number of intermediate tracks. At step 833, a question is asked as to whether there are any non-black events in the source track. A non-black event contains audio material that is currently activated for playback on a particular track. If answered in the negative, control is directed to step 837. Alternatively, control is directed to step 834. At step 834 the next non-black

25 event in the track is selected. At step 835 a question is asked as to whether it is possible to move the selected event to the destination track. If there is sufficient room on the destination track for the entire length of the selected event, then it is possible to move the event. If the question asked at step 835 is answered in the negative, control is directed to step 833, where

another event is considered for moving. Alternatively, control is directed to step 836, where the selected event is moved to the destination track. Thereafter, control is directed back to step 833, where any remaining non-black events on the source track are considered.

5 If the question asked at step 833 is answered in the negative, and no suitable events remain on the source track, control is directed to step 837. At step 837 the destination track is incremented. At step 838 a comparison is made between the value of the destination track and the value of the source track. If the destination track is less than the source track, control is
10 directed to step 833, where the next destination track can be considered. Alternatively, control is directed to step 839, where the source track is incremented. Thereafter, at step 840, a comparison is made between the value of the source track and the number of intermediate tracks. If these values match, this condition indicates that all the tracks have been
15 considered, and the optimisation process is complete. Alternatively, control is directed back to step 832, where each track in turn, below the value of the source track, is considered as a potential destination track for any events that remain on higher valued tracks. Numerically, the track optimisation proceeds as follows. On the first pass, track one is considered
20 as the source, with only track zero being considered as a destination. On the next iteration, track two is considered as the source, and tracks zero and one are considered as destinations. Next, track three is considered as a source, with tracks zero, one and two considered as destinations. This process is continued for as many source tracks as there are intermediate
25 tracks allocated at step 802. In considering destination tracks in this order, of the lowest going up to the highest, it is most likely to place events on lower numbered tracks. In this way, events are moved, whenever possible, up to lower numbered tracks, and in many practical situations this will result in a number of higher numbered tracks being completely empty at the end

of the optimisation process. Furthermore, the algorithm shown in *Figure 8B* operates in such a way that the empty tracks will be contiguous. In other words, it will be impossible, at the end of the process, for track four to be empty, and track five to contain events. Instead, there will always be contiguous blocks, so that, for example, tracks zero to four will contain events, and tracks five to twelve will be completely empty. This condition is useful, as it simplifies the process of deleting empty intermediate tracks, as indicated at step 808 shown in *Figure 8A*.

The effect of the optimisation procedures shown in *Figure 8A* and *Figure 8B*, when applied to the condition shown in *Figure 3*, is detailed in *Figure 9*. Dissolve 354 would be split up at step 805 and at step 807 optimisation would be performed. A gap larger than the length of the dissolve 354 is available in audio track A2, therefore the head of incoming track 356, required for frames 313 to 320, is transferred to audio track 2. Thus, by placing the head of this track in existing audio track 2, it is not necessary to use audio track 7 and all of the required audio data can be replayed in real time through the primary active tracks A1 to A6.

Attribute data for the transferred track is still derived from the attribute track associated with audio track A1. The head of the incoming track has been transferred, and therefore the attribute data available in audio track A1 for frame periods 313 to 320 relates to the tail of outgoing track 355. From frame 321 onwards, attribute data becomes available for the material of source 354. Attribute data for the head of source 354, the audio information which has been transferred to audio track A2 in accordance with the procedures shown in *Figure 9*, is derived by interpolating back from values specified from frame 321 onwards. Thus, this allows volume and pan attributes to be calculated for the head of source 354, which are in turn processed with the fading attributes of the dissolve; in order to achieve the required ramping-up of the audio levels as the head

of source 356 is played from audio track A2. In all other respects the audio information defined at portion 801 is replayed in the same way as other audio information in accordance with procedures detailed in *Figure 10*.

Procedures 714, for mixing audio data read from an audio buffer, are detailed in *Figure 10*. In the present embodiment, attributes are defined for each video frame, therefore at step 1001, where the next frame is selected for mixing, attribute data for the frame under consideration is loaded. Output values for the left and right channels are now generated on a sample by sample basis so as to accumulate the required data within playback buffer 605. At step 1002 the left and right accumulator buffers for the duration of the frame are initialised to zero values. At step 1003 the first track is selected. At step 1004 the first sample of the frame is selected. At step 1005 the next left and right samples are read from the track. At step 1006 the left and right output sample values are determined, in response to frame attribute data. At step 1007, these new values are added to the left and right accumulators for that sample number within the frame. At step 1008, a question is asked as to whether any samples remain to be calculated for the current frame. If answered in the affirmative, control is directed back to step 1004, where the next left and right sample pair is considered. Alternatively, control is directed to step 1009, where a question is asked as to whether another track is to be considered for accumulation. If answered in the affirmative, control is directed back to step 1003, where the next track is considered, and steps 1004 to 1009 are repeated. Finally, once all tracks have been considered for the current frame, control is directed to step 1010, where the contents of the accumulator buffers are first clipped, and then copied to the final playback buffers. It can be appreciated that the procedures illustrated in *Figure 10* require significant processing overhead. Thus, by reducing the number of audio tracks which need to be processed in order to supply data to the playback buffer, a

significant advantage is provided in terms of releasing processor resource.

The present embodiment has been described on the basis that normal operation is permitted for a certain number of audio tracks whereafter satisfactory processing is not possible if this number of tracks is exceeded. In some applications, a predetermined value for the total number of real time audio tracks permitted may be unspecified and the processing capabilities of a system may be variable depending on the selection of particular requirements. Thus, a user may be able to configure the operation of a system so as to provide optimal operating characteristics for a particular application. Under these circumstances, it may be preferable to optimise the presence of audio dissolves even when relatively few audio tracks have been enabled, on the basis that the number of active audio tracks should always be minimised so as to make processing capabilities available for use on other functions. In this way, further functions may be added to the system and made available when resources permit.

The present invention has been described with respect to the processing of stereo audio tracks. The invention may also be applied to the processing of video or other media tracks. In particular, given that tracks of any media type are represented and manipulated symbolically within the invention, no significant bandwidth restrictions are encountered when applying the invention to media of widely differing types, including high resolution video or film data. Thus, in systems where the number of tracks is limited by processing bandwidth, the invention may be applied advantageously. For example, there may be four or more audio channels per audio track, as is required for surround sound or multi-lingual sound systems. The present embodiment does not automatically result in an optimum arrangement of intermediate audio tracks, although it will usually result in an improvement, and never a degradation in overall system performance. In an alternative embodiment, the process 807 of optimising

intermediate tracks, includes steps to obtain an optimal solution. In order to identify an optimal solution, movement of events is prioritised according to their duration. Thus, events which have the longest duration have the higher priority with respect to track movement. In this way, short duration events, which are less likely to overlap, will congregate on tracks with a higher number. A number of passes through this algorithm will enable the optimal solution to emerge, and should not usually result in an excessive extra amount of processing to be performed. The optimal solution is characterised by the condition that the total number of tracks after optimisation is equal to the maximum number of tracks active at any one time, and this is a suitable test for the end condition of the alternative optimisation method.

Claims

1. Editing apparatus, including storage means configured to store digital samples;
5 display means configured to display symbolic representations of edit events within tracks; and
processing means configured to identify event locations and to move portions of edit events to alternative tracks so as to enhance processing performance.
10
2. Apparatus according to claim 1, wherein said processing means is configured to identify an event defined on a single track but requiring two interacting tracks and to transfer material from one of said two interacting tracks to a blank region of another track, thereby allowing both
15 of said interacting tracks to be played without allocating additional track resources.
3. Apparatus according to claim 1 or claim 2, wherein said processing means is configured to identify additional tracks; identify
20 overlapping events on the same track; move an overlapping event to a different track; identify improved locations for events; and identify a reduced number of tracks for subsequent processing.
4. Apparatus according to any of claims 1 to 3, wherein events
25 processed by said processing means are associated with collections of samples of audio data stored by said storage means.
5. Apparatus according to claim 3, wherein said overlapping events detected by said processing means define an audio cross-fade.

6. Apparatus according to claim 3, wherein said display means is configured to display an overlap of audio material in a form substantially similar to a video wipe or a video dissolve.

5

7. Apparatus according to claim 4, wherein said processing means is configured to process audio samples read from said storage means in combination with digital video samples read from said storage means.

10

8. Apparatus according to claim 7, wherein said display means is configured to display audio tracks as timelines presented against a shared time axis.

15

9. Apparatus according to claim 8, wherein the head of an event is moved and parameters for the moved event are determined with respect to parameters for the remaining material.

20

10. Apparatus according to claim 8, wherein said storage means is configured to record audio parameters for each video frame's-worth of audio samples.

25

11. A method of processing audio-visual data in which digital samples are read from storage media in response to edit events identified symbolically within tracks, wherein

an improved event location is identified on a different track, and said identified event is moved to said improved location so as to reduce the overall processing requirement.

12. A method according to claim 11, wherein an event defined on a single track but requiring two interacting tracks is identified; and

material for one of said two interacting tracks is transferred to a blank region of another track, thereby allowing both of said interacting tracks to be played without allocating additional track resources.

13. A method according to claim 11 or claim 12, including the steps of:

identifying additional tracks; identifying overlapping events on the same track; moving an overlapping event to a different track; identifying improved locations for events, and identifying a reduced number of tracks for subsequent processing.

14. A method according to any of claims 11 to 13, wherein said events are associated with collections of samples of audio data.

15. A method according to claim 13, wherein said overlapping events form an audio cross-fade.

16. A method according to claim 13, wherein said overlap is represented in a form substantially similar to a video wipe or a video dissolve.

17. A method according to claim 14, wherein said audio samples are processed in combination with digital video samples.

18. A method according to claim 17, wherein video tracks are displayed with audio tracks as timelines presented against a shared time axis.

19. A method according to claim 18, wherein the head of an event is moved and parameters for the moved event are determined with respect to parameters for the remaining material.

5

20. A method according to claim 18, wherein audio parameters are recorded for each video frame's-worth of audio samples.

21. A computer system loaded with executable instructions to perform the steps of displaying edit events symbolically in the form of a plurality of tracks;

10

identifying event locations;

moving portions of edit events to an alternative track to enhance processing performance; and

15

reading digital samples from a storage device in response to symbolic representations.

20

22. A computer system according to claim 21, wherein an event defined on a single track but requiring two interacting tracks is identified; and material for one of said interacting tracks is transferred to a blank region of another track, thereby allowing both of said interacting tracks to be played without allocating additional track resources.

25

23. A computer system according to claim 21 or 22, wherein said executable instructions are configured such that said system identifies additional tracks, identifies overlapping events on the same track, moves an overlapping event to a different track, identifies improved locations for events and identifies a reduced number of tracks for subsequent processing.

24. A computer system according to any of claims 21 to 23, programmed such that said events are associated with collections of samples of audio data.

5 25. A computer system according to claim 23, programmed such that overlapping events representing an audio cross-fade are detected.

26. A computer system according to claim 23, programmed such that said audio overlap is represented in a form substantially similar to a
10 video wipe or a video dissolve.

27. A computer system according to claim 24, programmed to process said audio sample in combination with digital video samples.

15 28. A computer system according to claim 27, programmed to display video tracks with audio tracks as time lines presented against a shared time axis.

29. A computer system according to claim 28, programmed to
20 move the head of an event and parameters for the moved event with respect to the parameters for the remaining material.

30. A computer system according to claim 28, programmed to record audio parameters for each video frame's-worth of audio samples.

25 31. A computer-readable medium having computer-readable instructions executable by a computer such that said computer performs the steps of:

displaying edit events symbolically in the form of a plurality of tracks;

identifying event locations;

moving portions of edit events to an alternative track to enhance processing performance; and

reading digital samples from a storage device in response to said symbolic representations.

32. A computer-readable medium according to claim 31, having computer-readable instructions executable by a computer such that said computer performs the further step of transferring material for one of said interacting tracks to a blank region of another track, thereby allowing both of said interacting tracks to be played without allocating additional track resource.

33. A computer-readable medium according to claim 31 or 32, having computer-readable instructions executable by a computer such that said computer performs the further step of identifying additional tracks, identifying overlapping events on the same track, moving an overlapping event to a different track, identifying improved locations for events, and identifying a reduced number of tracks for subsequent processing.

34. A computer-readable medium according to any of claims 31 to 33, having computer-readable instructions executable by a computer such that said computer performs the further step of associating said events with collections of samples of audio data.

35. A computer-readable medium according to claim 33, having computer-readable instructions executable by a computer such that said computer performs the further step of associating overlapping events that define an audio cross-fade.

5 36. A computer-readable medium according to claim 33, having computer-readable instructions executable by a computer such that said computer performs the further step of representing an audio overlap in a form substantially similar to a video wipe or a video dissolve.

10 37. A computer-readable medium according to claim 34, having computer-readable instructions executable by a computer such that said computer performs the further step of processing audio samples in combination with digital video samples.

15 38. A computer-readable medium according to claim 37, having computer-readable instructions executable by a computer such that said computer performs the further step of displaying video tracks with audio tracks as timelines presented against a shared time axis.

20 39. A computer-readable medium according to claim 38, having computer-readable instructions executable by a computer such that said computer performs the further step of moving the head of an event and parameters for the moved event with respect to parameters for the remaining material.

25 40. A computer-readable medium according to claim 38, having computer-readable instructions executable by a computer such that said computer performs the further step of recording audio parameters for each video frame's-worth of audio samples.

41. A computer-readable memory system having a plurality of data fields stored therein representing a data structure, wherein said structure include

5 symbolic representations of edit events referenced as belonging to specified tracks, wherein portions of edit data are relocatable to enhance playback capabilities.

42. A computer-readable memory system having data fields stored therein according to claim 41, wherein said structure further includes
10 two interacting tracks represented as a single track, wherein material for one of said interacting tracks is transferred to a blank region of another track, thereby allowing both of said interacting tracks to be played without allocating addition track resources.

43. Editing apparatus substantially as herein described with
15 reference to the accompanying drawings.

44. A method of processing audio-visual data substantially as herein described with reference to the accompanying drawings.



Application No: GB 9906881.9
Claims searched: 1 to 44

Examiner: Donal Grace
Date of search: 12 July 1999

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK Cl (Ed.Q): G5R (RB81)
Int Cl (Ed.6): G11B 27/028, 27/034
Other: Online: EPODOC; JAPIO; WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X, A	GB 2289558 A (AVID TECHNOLOGY)	X: 41, 42 A: 1
X, A	US 5726717 (PETERS et al)	X: 41, 42 A: 1

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☒ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.